FALL 2006
COMPUTER SCIENCES DEPARTMENT
UNIVERSITY OF WISCONSIN-MADISON
PH. D. QUALIFYING EXAMINATION
Networking
Monday, September 18, 2006
3:00-7:00 PM

GENERAL INSTRUCTIONS:

1. Answer each question in a separate book.

2. Indicate on the cover of *each* book the area of the exam, your code number, and the question answered in that book. On one of your books list the numbers of all the questions answered. *Do not write your name on any answer book.*

3. Return all answer books in the folder provided. Additional answer books are available if needed.

SPECIFIC INSTRUCTIONS:

All six questions must be answered.

POLICY ON MISPRINTS AND AMBIGUITIES:

The Exam Committee tries to proofread the exam as carefully as possible. Nevertheless, the exam sometimes contains misprints and ambiguities. If you are convinced a problem has been stated incorrectly, mention this to the proctor. If necessary, the proctor can contact a representative of the area to resolve problems during the *first hour* of the exam. In any case, you should indicate your interpretation of the problem in your written answer. Your interpretation should be such that the problem is nontrivial.

Computer Sciences Department
Networking Qualifying Exam
Spring 2006
All six questions must be answered.

1. **New Internet Architectures**

The architecture of today's Internet has been guided over the years by a number of simple principles such as a single IP network protocol and end-to-end arguments. However, there are inherent limitations in the current design that have led some to think about redesigning the basic architecture.

   a. One new design concept that is currently being considered is "network virtualization". The idea is to allow multiple, different networks (i.e. multiple versions of an Internet Protocol) to exist at layer 3 the same time. What are the implications both positive and negative about this idea?
   b. It can be argued that end-to-end arguments have shaped much of the design and implementation of today's Internet. Describe the characteristics of today's Internet that result from end-to-end arguments and then describe how the arguments might be modified to address a specific example of a problem in today's Internet (e.g. spam, malicious attacks, or network management).

2. **Issues in Congestion Control**

Congestion is one of the basic risks associated with statistically multiplexed packet switched networks. Congestion is typically managed by end hosts using some variant of Jacobson's original TCP-Tahoe. While these methods have enabled the Internet to scale to it's current size, there are situations in which current versions of TCP can be quite inefficient.

   a. Consider a network with extremely high bandwidth links with long paths from end to end, and there is a need to transfer very large files. Describe at least two potential problems with using an out-of-the-box version of TCP-Reno and how can one take better advantage of high bandwidth paths?
   b. Optical switching devices enable extremely high bandwidth virtual circuits between end points to be established at relatively course timescales (e.g. 1ms). Their use, however, has been limited by a number of issues including the fact that they can only buffer a small number of bytes. Explain why this is a problem in operational networks and suggest how this problem might be addressed.

3. **Wireless networking**

Wireless Mesh Networks (WMNs) claim to be a viable mechanism to provide Internet access services in many urban scenarios. Such a network consists of a number of wireless routers (called mesh points) that provide multi-hop wireless connectivity between clients and wired gateways. The wireless nature of all links in WMNs imply that mechanisms in wired networks are often not applicable.

   a. Explain why a standard link state routing protocol, such as OSPF, designed for wired networks, may not perform well in WMNs.
   b. Performance of TCP on a multi-hop wireless path is worse than its performance on a multi-hop wired path. Why?
   c. The 802.11 MAC protocol was specifically designed for a single hop wireless environment and works best when all wireless nodes are in range of each other. Explain why this protocol is not adequate for multi-hop wireless environments.

4. **Internet services**

As the Internet has grown in popularity and scale, it has become increasingly difficult to experiment with and deploy new services at the lower layers of the protocol stack. Therefore, overlay networks emerged as a mechanism to overcome this limitation in deploying new services.

a. Describe two advantages of using overlay-based mechanisms in the deployment of a multicast service.
b. In many cases, Network Address Translators (NATs) adversely affect deployment of specific overlay-based services. Explain how.
c. Can a Quality of Service mechanism be deployed using an overlay-based approach? Explain your answer.

5. **Security and DoS**

Denial-of-Service (DoS) attacks are a major menace to public Internet servers. Several mechanisms have been proposed in the literature to prevent or respond to such attacks. However, there are few working solutions in place today.

a. A common tactic to evade identification of the source(s) of attacks is to spoof the source IP on attack packets. Ingress filtering, where an ISP checks if a packet entering its network has a legitimate source address, has been proposed to curtail source spoofing. Outline the challenges involved in getting ingress filtering to effectively filter spoofed packets on the Internet.
b. SYN cookies is another popular mechanism to offer protection against a class of DoS attacks. The SYN cookie mechanism works as follows: To open a TCP connection with the server, a client first sends a TCP SYN. In response, the server sends a TCP SYN+ACK packet back to the client. Like all TCP packets, the server's SYN+ACK also carries a "sequence number". The trick in SYN cookies is that the server computes a special initial sequence number as follows:
   - First k bits: C mod 32, where C is a counter incremented every 64s.
   - Remaining (32-k) bits: a one way hash function computed over the connection 4 tuple, and C. The function is known only to the server.
   1. Assume that clients use regular, unmodified TCP stacks. Outline the actions taken by a client when it receives the above SYN+ACK.
   2. What forms of resource exhaustion attacks can a server using SYN cookies successfully avoid? Explain your answer.
   3. Outline two attacks against which SYN cookies are ineffective. Explain your answer.

6. **Routing**

After taking CS740, Hugh Hopeful proposed a novel routing/forwarding solution that could replace IGPs based on shortest path routing and better balance the traffic within the ISP's network. The core idea behind his proposal is that routers need not send traffic to the next hop that is on the shortest path to the destination, but instead they can split the traffic between all next hops that are closer to the destination (next hops whose shortest distance to the destination is smaller than the shortest distance of the current router). For all such possible next hops, the traffic would be split randomly with probability proportional to the free capacity on the links towards those next hops. For example, if there are two feasible next hop routers for the destination of a given packet and the link to one of them is 80% loaded (20% free) while the link to the other one is 40% loaded (60% free) with probability 25% the packet will go towards the first next hop, and with probability 75% towards the second next hop.

a. Does this new routing/forwarding solution guarantee that after the routing system stabilizes, packets will not loop in the network? If no, give a counterexample, if yes, sketch a proof.
b. Hugh's scheme is used by a network with the topology below. The traffic mix consists of 4 constant bit rate UDP streams of 6 Mbps of traffic each from S1 to D1, from S2 to D1, from S3 to D2 and from S4 to D3. Assume that all links have a cost of 1 in the routing protocol, they can carry 12Mbps, and if the offered load is larger, traffic is dropped randomly. Indicate for each of the 4 streams of traffic what the end to end loss rate will be. Can the overall amount of dropped traffic be reduced?

```
S1---R1                 D1
       \               /
        \             /
S2--R2--R4-------R6
         /           \
```

```
        /           \
S3---R3              D2
     / \            /
   /    \          /
  /     R5-------R7
S4                 \
                    \
                     D3
```

c. Hugh's proposal can reorder packets belonging to the same flow by sending them on paths with different propagation delays. Propose a change that will maintain the benefits of load balancing, but reduce the reordering problem. Your solution should not require routers to keep per flow state. You can assume that each flow represents a small percentage of the traffic towards a destination and that the overall traffic is relatively stable.